

Experiences with a Distributed Virtual Patient Record System

David W. Forslund, PhD, Richard L. Phillips PhD, David G. Kilman, MS, James L. Cook, MD*

Los Alamos National Laboratory

Los Alamos, New Mexico

*National Jewish Center for Immunology and Respiratory Medicine

Denver, Colorado

TeleMed is a distributed diagnosis and analysis system, which permits physicians who are not collocated to consult on the status of a patient. The patient's record is dynamically constructed from data that may reside at several sites but which can be quickly assembled for viewing by pointing to the patient's name. Then, a graphical patient record appears, through which consulting physicians can retrieve textual and radiographic data with a single mouse click. TeleMed uses modern distributed object technology and emerging telecollaboration tools.

INTRODUCTION

Healthcare is undergoing a revolution because of technological, political and sociological changes. In order for these changes to result in cost effective, quality healthcare, a fundamental change in the way healthcare providers and payers manage patient information is needed throughout the nation and around the world. The change requires treating patient data in a completely different way than has been considered until recently.¹ In particular, it is useful to think of the patient's medical record as an entity that exists on the network in its aggregate form simultaneously populated from multiple locations. We describe in this paper some of the motivation for this change, what we mean by a virtual patient record, and some results of some early implementations of a virtual patient record.

BACKGROUND

Information technologies have begun to have profound impact on a variety of business and social applications. These technologies both improve the quality and lower the costs of business processes. Healthcare is no different in this respect. As an example, the increased mobility of the patient populations and changes in healthcare providers and payers has resulted in a patient's medical information being accumulated in a variety of locations -- hospitals, HMO's and doctor's offices -- with little or no linkage between them. Because of these multiple points of entry of patient information into the healthcare system, both healthcare provider and payer get a fragmented picture of the health history of a patient, particularly if she has some kind of chronic illness such as asthma or diabetes. This fragmented view can occur over a regional network of clinics as well as over the entire country. We call this requirement for multiple entry points into the healthcare system "distributed healthcare." Because of distributed healthcare, the patient frequently is the individual

with the most complete historical information as to how their clinical illness has progressed.

In addition to the trend towards distributed healthcare, there is a rapid movement to computerized patient records within hospitals and HMO's. However, even electronic access over a region may not be sufficient to track a significant number of patients as job mobility increases. This increasingly wide diffusion of the population requires patient data to be accessible in an organized manner on a national and even global scale, independent of the healthcare provider or payer. To deal with this information explosion, there are a number of organizations working to standardize healthcare information and communication including the JWG-CDM (Joint Working Group for a Common Data Model), the Health Level 7 (HL7) group, the CPRI (Computerized Patient Record Institute) and recently the Object Management Group (OMG).²

The impact on the healthcare industry of making healthcare information available over wide areas in a secure manner will be quite profound. Such availability could potentially allow for data mining of information. This information could then be used to discover and analyze associations between disease entities and previously unknown risk factors (recorded in the patient history), to test hypotheses regarding putative risk factors, or to study disease distribution using demographic data. Applications of data mining could also include enabling a physician to do a comparative analysis of a particular patient's symptoms with the symptoms of other patients with similar or different diseases. Having wide area access to healthcare information would allow for more intelligent video consultations. During these consultations, along with the video, specialists in multiple locations could simultaneously see and annotate a patient's record. HMO's could do a better job of outcomes analysis, physicians would have access to better decision support information, and patients could be better educated to manage their health. All of these applications require advanced pattern matching techniques beyond simple database searches.

VIRTUAL PATIENT RECORD CONCEPT

For a long time healthcare providers and payers have realized that electronic records have real value and that moving this digital data around to areas where it is needed is highly desirable.³ There has been considerable success at the Veterans Affairs Hospitals

and other locations implementing a decentralized patient record system, but these systems have not addressed the issue of doing so over a wide area network between different domains.⁴ With legacy systems, moving data between disparate databases has been a real problem that has been well addressed by the HL7 effort. However, we believe that for ease of access by end users such as physicians and patients, the patient's information must appear to the user as a unified set of data even though it may be spread all over the country. The user's view, of course, might access only a specially tailored subset of the records in order to handle issues of displaying the information in an intelligible manner. With distributed object technology, which can hide much of the vagaries of accessing information, such a view of data is now possible.⁵

The virtual patient record is virtual in that it is a view of the data that might be configured differently at different locations, but that is mapped into a common format at the time the record is required. Creation of the virtual patient record must be done with minimal compromise in the integrity of the data while maintaining high accessibility. For example, simple store and forward systems have potentially serious difficulties because data is copied to multiple locations and then edited and amended locally. There typically is no mechanism to integrate new information entered into any local copy back into the primary record and all other local copies without considerable human effort.

Through a virtual patient record distributed healthcare data is made available through references (analogous to hypertext links of the World Wide Web) and is only brought together (or created) on demand by the end user. Since users generally access components of a record rather than the entire patient record, data movement is minimized. In the distributed system, reference counting capabilities and distributed transaction processing maintain the integrity of the data. Thus, full asynchronous access of the record that enables multiple physicians, other healthcare providers and healthcare payers to update the patient record is supported. Many of these capabilities have already been specified, for example, in the new CORBA services developed by the OMG.

This model does require ubiquitous network connectivity and accessibility, but high bandwidth transmission is not necessary unless large amounts of image or video data need to be moved. Such an infrastructure is rapidly coming into existence, even in rural areas. The model also requires a robust security infrastructure to support authentication, confidentiality, and data integrity so that there is no single point of failure that, if compromised, would give access to all the information.⁶ In order to provide robust data access even when larger numbers of users are attempting to access data and a universal but secure way to identify and locate patient information, various replication servers are also required.

We believe that it is now possible to implement the virtual patient record concept if all stakeholders -- government, public, and private -- cooperate in making it a reality in the everyday practice of medicine. Many of the underlying standards are being put into place, but more standard representations of medical objects are needed. This is one of the goals of the new Healthcare Task Force created by the OMG (and also known as CORBAmed).

Viability of the virtual patient record will depend heavily on the ability to quickly and securely identify patients and their respective healthcare providers and payers. This requirement can be met by a Master Patient Index (MPI). Besides the basic architecture to enable healthcare objects to work interchangeably and together, in order to avoid the chaos caused by using existing naming conventions, the virtual patient record system will have to be adopted by a large portion of the healthcare community in a short period of time.

TELEMED: A PROTOTYPE EXPERIENCE

In a joint effort with physicians at the National Jewish Center for Immunology and Respiratory Medicine (NJC), we have constructed a prototype of such a system described above. This TeleMed system enables physicians who are not collocated to simultaneously see, edit and annotate a patient record at remote locations. It handles multimedia data including CT imaging and audio annotations. It uses Object Request Brokers (ORBs) which abstract the distributed databases that provide the persistent object storage of the multimedia data. It has object-level security implemented to provide authentication and encryption for confidentiality. It is built with the idea of providing easy-to-use access to complex information while providing advanced data mining techniques accessible to an end user. TeleMed is, therefore, an early implementation of the virtual patient record described above and demonstrates that the concept is achievable. It has been deployed at the NJC, the National Institutes of Health, and at the Texas Medical Center for early testing and evaluation. Physicians at these three institutions can simultaneously view, edit, and annotate the patient data stored at multiple locations while each physician can see the data the other physician has entered. To the physician using TeleMed, it appears as if all the data resides on their own desktop computer; there is no indication that multiple databases are involved. We have also implemented, where the available bandwidth permits, the ability to support video teleconferencing within the TeleMed system.

TELEMED OVERVIEW

Some of the capabilities of TeleMed can be illustrated by looking at a series of user interface components that are available to the user. The user begins a TeleMed session by selecting a database site from the interface shown in Figure 1. This sets in motion a CORBA-based transaction for vending all patient record objects from the selected site to the requesting client, shown listed in Figure 1.

To understand the coordination of distributed object activities with user interface activities we can consult Figure 2, a graphical representation of TeleMed objects. In this diagram the arrows represent an inheritance relationship and the other lines represent a reference or containment. Textual data from the Patient object, i.e., the patient's name, was retrieved and used to populate the patient list in Figure 1.



Figure 1 - Initial TeleMed Interface

A patient's treatment record appears by double-clicking on the patient's name in the interface in Figure 1. The user interface manifestation of the Observation Battery, Treatment, and Annotation objects in Figure 2 is called a Graphical Patient Record (GPR) and is shown in Figure 3.

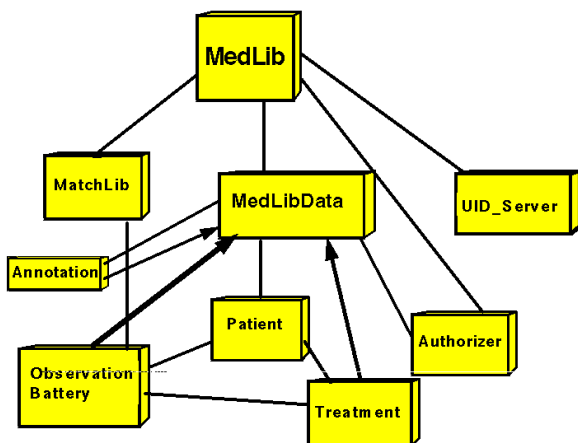


Figure 2 - TeleMed Objects

The GPR is an excellent example of media-rich document and distributed object technology. The GPR is a virtual document, a patient record that is empty until it is dynamically populated by requests for distributed objects. The Observation Battery, etc. objects in Figure 2 contain the information necessary

for "harvesting" the data from all appropriate sites. Thus, laboratory reports may be retrieved from the National Institutes of Health in Bethesda, MD while radiographic data may reside at the NJC in Denver. So, for example, when all patient data are retrieved, icons representing laboratory tests, radiographic studies, drug treatments, etc. are drawn on the GPR template.

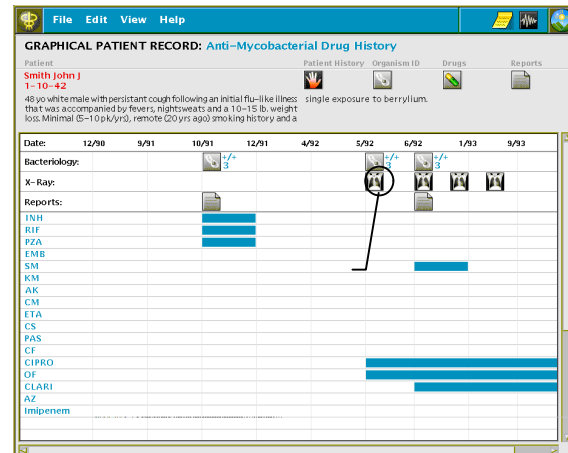


Figure 3 - TeleMed Graphical Patient Record

Each of these icons is mouse-sensitive and, when clicked, summon up additional user interfaces and related data.

Before looking at these interfaces it will be helpful to know more about what goes on at the distributed objects level. In Figure 4 we show the relationships between the client process (TeleMed GUI) and the two controlling objects, MedLib and MatchLib. Any of these three entities can reside at any location. In fact, the TeleMed GUI can communicate with any number of MedLib objects, which, in turn, can call upon the services of any number of MatchLib objects. Suppose, now, the user clicks on a CT study icon in the GPR in Figure 3. This causes a request to be sent to the current MedLib to retrieve that patient's CT study from the corresponding persistent object store. That transaction causes the user interface shown in Figure 5 to appear.

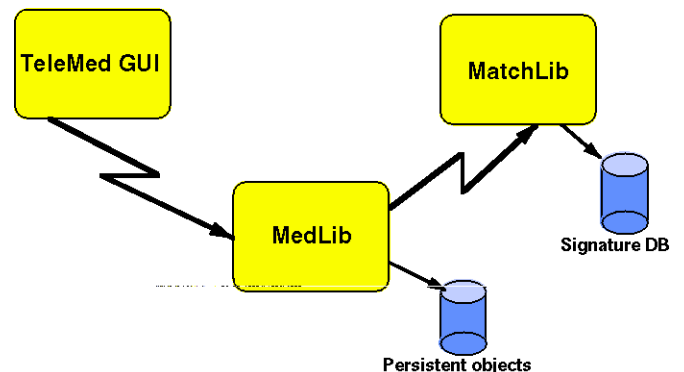


Figure 4 - TeleMed Object Distribution

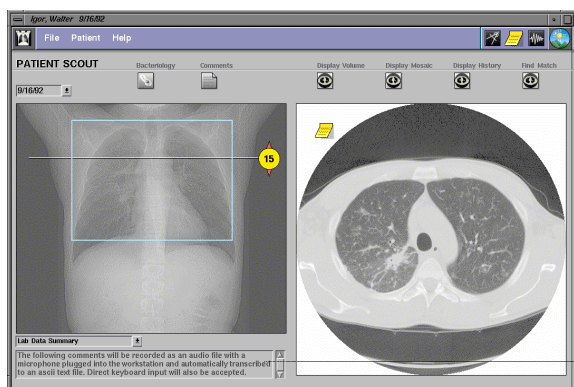


Figure 5 - TeleMed CT Study Interface

The image on the left of Figure 5 is a scout, so named because it was originally used by the CT technician as a guide in determining where to produce full transverse slice images of the patient. In this interface the scout is similarly used, but now as a guide for the physician in selecting slices to view from the database. To do this, the horizontal cursor is dragged up or down to the desired location and released. Here, slice number 15 was selected and is shown on the right.

SECURITY

Protecting patient confidentiality is the primary security concern for any medical data. Also important are protections against unauthorized additions, deletions and modifications of patient data. The security infrastructure layer of TeleMed is intended to provide the security services necessary to allow access control for patient data. This includes two fundamental services: an authentication framework and secure remote method calls. Any particular access control policy is implemented at a higher level, using an application-specific authorization object and access control lists.

Security is introduced in the system through filters and transformers in the method stubs in the GUI that interact with the ORBs.

A key and ticket server pair provide fundamental authentication services for both human users and CORBA server objects, and also provide a mechanism for secure session key exchange. The authentication and key exchange protocol is similar to Kerberos. Users register RSA public/private key pairs with the key server, while DES secret keys are registered for CORBA objects.

Encrypted RSA private keys are stored in the key server to allow convenient remote access by users. By retrieving private keys via CORBA, users need not have their keys stored on a local file system. Private keys are encrypted with a DES secret key (constructed from a user-selected pass-phrase) before

submission to the key server, so the private keys are never disclosed to a system administrator.

Like Kerberos, the authentication mechanism is implicit. Users authenticate their identity by virtue of their ability to successfully decipher the encrypted session key. Servers authenticate their identity by virtue of their ability to successfully decipher the encrypted ticket. The ticket server authenticates its identity by signing the session key and ticket it produces (using an RSA private key).

It is worth noting that a client may be a CORBA server acting on behalf of another client of its own. For example, in the TeleMed system, when a user requests a merged patient record, a MedLib server must contact other MedLib servers to gather references to the distributed patient data. In such a case, the object must be able to establish a connection with exactly the privileges appropriate to its own client. That is, in some cases, a server object must be able to forward the credentials of its client to other servers.

We plan on evolving this design to be in compliance with the CORBA security model that has been recently adopted by the OMG.

ADVANCED CAPABILITIES

TeleMed supports basic data mining by providing abilities to compare images with “similar” features and to visually navigate through a image data base. We conclude our discussion of the TeleMed application by describing one of its most powerful features. This feature allows a user to perform a “query by example” search of an image database. Many technologies are represented in this feature — massively parallel computation servers, image analysis agents, and distributed object computing. To be specific, the MatchLib object shown in Figure 4 encapsulates the image analysis agent as a member function. For best performance MatchLib will typically reside on a massively parallel computer because the matching algorithm is inherently parallel. The signature database, which contains representative features of each image, usually resides on the same machine as MatchLib. Finally, the user invokes this entire matching operation simply by clicking the “Find Match” button in the upper right of Figure 5. The selected slice is used as the query image. The result of a matching operation is shown in Figure 6.

In Figure 6 the upper left image is the same one the user specified as the example query image. The result of the match is summarized by the thumbnails in the lower scrolling window. Clicking in a selected thumbnail causes its full-size representation to appear in the upper right comparative inspection window.

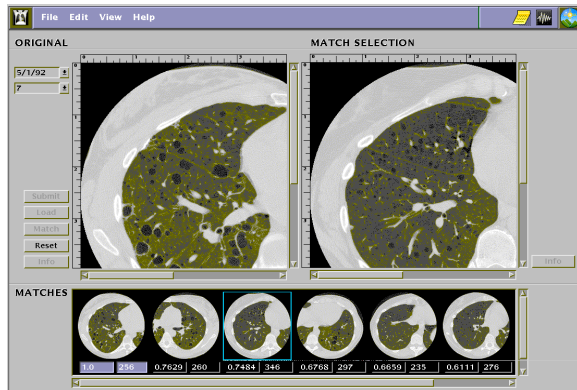


Figure 6 - TeleMed Image Matching Interface

CONCLUSION

One virtue of the virtual patient record system is that it can be extended to a wide variety of clinical and diagnostic areas. We believe that the cooperation of all stakeholders working with together to build a common infrastructure will not only develop new business opportunities but also will make a positive impact on the healthcare delivery nationally and worldwide. The opportunity exists now to invest in a new healthcare infrastructure that will significantly enhance the delivery of quality of healthcare at a reasonable cost.

References

1. Phillips RL, A Network-based Distributed, Media-rich Computing and Information Environment, Digital Media and Electronic Publishing conference sponsored by the British Computer Society in Leeds, UK, December 1994.
2. Object Management Group, Ed.. *Common Object Request Broker Architecture* Wiley-QED, 1991.
3. Dick R and Steen E, eds. *The Computer-Based Patient Record, An Essential Technology for Health Care*, Institute of Medicine, National Academy Press, 1991.
4. Decentralized Hospital Computer Program, Department of Veteran Affairs, August, 1995. See also: <http://www.va.gov/vama.htm>
5. Orfali R, Harkey D, Edwards J, *The Essential Distributed Objects Survival Guide*, John Wiley & Sons, Inc. 1996.
6. Donaldson M and Lohr K, eds. *Health Data in the Information Age: Use, Disclosure, and Privacy*, Institute of Medicine, National Academy Press, 1994.